

# Optimal Listing of Cycles and $st$ -Paths in Undirected Graphs

R. Ferreira (U. Pisa, [ferreira@di.unipi.it](mailto:ferreira@di.unipi.it))      R. Grossi (U. Pisa, [grossi@di.unipi.it](mailto:grossi@di.unipi.it))  
 A. Marino (U. Firenze, [a.marino@unifi.it](mailto:a.marino@unifi.it))      N. Pisanti (U. Pisa, [pisanti@di.unipi.it](mailto:pisanti@di.unipi.it))  
 R. Rizzi (U. Verona, [r.rizzi@univr.it](mailto:r.rizzi@univr.it))      G. Sacomoto (INRIA, [sacomoto@gmail.com](mailto:sacomoto@gmail.com))

## Abstract

Listing all the simple cycles (hereafter just called cycles) in a graph is a classical problem whose efficient solutions date back to the early 70s. For a graph with  $n$  vertices and  $m$  edges, containing  $\eta$  cycles, the best known solution in the literature is given by Johnson's algorithm [SIAM J. Computing, 1975] and takes  $O((\eta + 1)(m + n))$  time. This solution is surprisingly not optimal for undirected graphs: to the best of our knowledge, no theoretically faster solutions have been proposed in almost 40 years.

We present the first optimal solution to list all the cycles in an undirected graph  $G$ , improving the time bound of Johnson's algorithm by a factor that can be  $O(n^2)$ . Specifically, let  $\mathcal{C}(G)$  denote the set of all these cycles, and observe that  $|\mathcal{C}(G)| = \eta$ . For a cycle  $c \in \mathcal{C}(G)$ , let  $|c|$  denote the number of edges in  $c$ . Our algorithm requires  $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$  time and is asymptotically optimal: indeed,  $\Omega(m)$  time is necessarily required to read  $G$  as input, and  $\Omega(\sum_{c \in \mathcal{C}(G)} |c|)$  time is necessarily required to list the output.

We then describe an infinite family of dense graphs, in which each graph with  $n$  vertices and  $m$  edges contains  $\eta = \Omega(m)$  cycles  $c$  with  $|c| = O(1)$  edges. For each graph in this family, our algorithm requires  $\Theta(\eta + m) = \Theta(\eta)$  time, thus saving  $O(m + n) = O(n^2)$  time when compared to Johnson's algorithm. In general for any graph, since  $|c| \leq n$ , the cost of our algorithm never exceeds  $O(m + (\eta + 1)n)$  time.

We also present the first optimal solution to list all the simple paths from  $s$  to  $t$  (shortly,  $st$ -paths) in an undirected graph  $G$ . Let  $\mathcal{P}_{st}(G)$  denote the set of  $st$ -paths in  $G$  and, for an  $st$ -path  $\pi \in \mathcal{P}_{st}(G)$ , let  $|\pi|$  be the number of edges in  $\pi$ . Our algorithm lists all the  $st$ -paths in  $G$  optimally in  $O(m + \sum_{\pi \in \mathcal{P}_{st}(G)} |\pi|)$  time, observing that  $\Omega(\sum_{\pi \in \mathcal{P}_{st}(G)} |\pi|)$  time is necessarily required to list the output.

While the basic approach is simple (see binary partition in point 3), we use a number of non-trivial ideas to obtain our optimal algorithm for an undirected (connected) graph  $G$  as conceptually listed below.

1. Prove the following reduction. If there exists an optimal algorithm to list the  $st$ -paths in  $G$ , there exists an optimal algorithm to list the cycles in  $G$ . This relates  $\mathcal{C}(G)$  and  $\mathcal{P}_{st}(G)$  for some choices  $s, t$ .
2. Focus on listing the  $st$ -paths. Consider the decomposition of the graph into biconnected components (BCCs), thus forming a tree  $T$  where two BCCs are adjacent in  $T$  iff they share an articulation point. Exploit (and prove) the property that if  $s$  and  $t$  belong to distinct BCCs, then (i) there is a unique sequence  $B_{s,t}$  of adjacent BCCs in  $T$  through which each  $st$ -path must necessarily pass, and (ii) each  $st$ -path is the concatenation of paths connecting the articulation points of these BCCs in  $B_{s,t}$ .
3. Recursively list the  $st$ -paths in  $B_{s,t}$  using the classical binary partition (i.e. given an edge  $e$  in  $G$ , list all the cycles containing  $e$ , and then all the cycles not containing  $e$ ): now it suffices to work on the *first* BCC in  $B_{s,t}$ , and efficiently maintain it when deleting an edge  $e$ , as required by the binary partition.
4. Use a notion of *certificate* to avoid recursive calls (in the binary partition) that do not list new  $st$ -paths. This certificate is maintained dynamically as a data structure representing the first BCC in  $B_{s,t}$ , which guarantees that there exists at least one *new* solution in the current  $B_{s,t}$ .
5. Consider the binary recursion tree corresponding to the binary partition. Divide this tree into *spines*: a spine corresponds to the recursive calls generated by the edges  $e$  belonging to the same adjacency list in  $B_{s,t}$ . The amortized cost for each listed  $st$ -path  $\pi$  is  $O(|\pi|)$  if there is a guarantee that the amortized cost in each spine  $S$  is  $O(\mu)$ , where  $\mu$  is a lower bound on the number of  $st$ -paths that will be listed from the recursive calls belonging to  $S$ . The (unknown) parameter  $\mu$  different for each spine  $S$ , and the corresponding cost  $O(\mu)$ , will drive the design of the proposed algorithms.

# 1 Introduction

Listing all the simple cycles (hereafter just called cycles) in a graph is a classical problem whose efficient solutions date back to the early 70s. For a graph with  $n$  vertices and  $m$  edges, containing  $\eta$  cycles, the best known solution in the literature is given by Johnson’s algorithm [7] and takes  $O((\eta + 1)(m + n))$  time. This solution is surprisingly not optimal for undirected graphs: to the best of our knowledge, no theoretically faster solutions have been proposed in almost 40 years.

**Our results.** We present the first optimal solution to list all the cycles in an undirected graph  $G$ , improving the time bound of Johnson’s algorithm by a factor that can be  $O(n^2)$ . Specifically, let  $\mathcal{C}(G)$  denote the set of all these cycles, and observe that  $|\mathcal{C}(G)| = \eta$ . For a cycle  $c \in \mathcal{C}(G)$ , let  $|c|$  denote the number of edges in  $c$ . Our algorithm requires  $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$  time and is asymptotically optimal: indeed,  $\Omega(m)$  time is necessarily required to read  $G$  as input, and  $\Omega(\sum_{c \in \mathcal{C}(G)} |c|)$  time is necessarily required to list the output. Since  $|c| \leq n$ , the cost of our algorithm never exceeds  $O(m + (\eta + 1)n)$  time.

Along the same lines, we also present the first optimal solution to list all the simple paths from  $s$  to  $t$  (shortly,  $st$ -paths) in an undirected graph  $G$ . Let  $\mathcal{P}_{st}(G)$  denote the set of  $st$ -paths in  $G$  and, for an  $st$ -path  $\pi \in \mathcal{P}_{st}(G)$ , let  $|\pi|$  be the number of edges in  $\pi$ . Our algorithm lists all the  $st$ -paths in  $G$  optimally in  $O(m + \sum_{\pi \in \mathcal{P}_{st}(G)} |\pi|)$  time, observing that  $\Omega(\sum_{\pi \in \mathcal{P}_{st}(G)} |\pi|)$  time is necessarily required to list the output. We prove the following reduction to relate  $\mathcal{C}(G)$  and  $\mathcal{P}_{st}(G)$  for some suitable choices of vertices  $s, t$ : If there exists an optimal algorithm to list the  $st$ -paths in  $G$ , then there exists an optimal algorithm to list the cycles in  $G$ . Hence, we can focus on listing  $st$ -paths.

**History of the problem.** The classical problem of listing all the cycles of a graph has been extensively studied for its many applications in several fields, ranging from the mechanical analysis of chemical structures [16] to the design and analysis of reliable communication networks, and the graph isomorphism problem [21]. There is a vast body of work, and the majority of the algorithms listing all the cycles can be divided into the following three classes (see [1, 11] for excellent surveys).

*Search space algorithms.* According to this approach, cycles are looked for in an appropriate search space. In the case of undirected graphs, the *cycle vector space* [3] turned out to be the most promising choice: from a basis for this space, all vectors are computed and it is tested whether they are a cycle. Since the algorithm introduced in [21], many algorithms have been proposed: however, the complexity of these algorithms turns out to be exponential in the dimension of the vector space, and thus in  $n$ . For the special case of planar graphs, in [17] the author was able to design an algorithm listing all the cycles in  $O((\eta + 1)n)$  time.

*Backtrack algorithms.* According to this approach, all paths are generated by backtrack and, for each path, it is tested whether it is a cycle. One of the first algorithms based on this approach is the one proposed in [20], which is however exponential in  $\eta$ . By adding a simple pruning strategy, this algorithm has been successively modified in [19]: it lists all the cycles in  $O(nm(\eta + 1))$  time. Further improvements were proposed in [7], [18], and [13], leading to  $O((\eta + 1)(m + n))$ -time algorithms that work for both directed and undirected graphs.

*Algorithms using the powers of the adjacency matrix.* This approach uses the so-called *variable adjacency matrix*, that is, the formal sum of edges joining two vertices. A non-zero element of the  $p$ -th power of this matrix is the sum of all walks of length  $p$ : hence, to compute all cycles, we compute the  $n$ th power of the variable adjacency matrix. This approach is not very efficient because of the non-simple walks. All algorithms based on this approach (e.g. [12] and [23]) basically differ only on the way they avoid to consider walks that are neither paths nor cycles.

Almost 40 years after Johnson’s algorithm [7], the problem of efficiently listing all cycles of a graph is still an active area of research (e.g. [2, 5, 6, 10, 14, 22, 15]). New application areas have emerged in the last decade, such as bioinformatics: for example, two algorithms for this problem

have been proposed in [8] and [9] while studying biological interaction graphs. Nevertheless, no significant improvement has been obtained from the theory standpoint: in particular, Johnson's algorithm is still the theoretically most efficient.

**Hard graphs for Johnson's algorithm.** We now describe an infinite family of dense undirected graphs, in which each graph with  $n$  vertices and  $m$  edges contains  $\eta = \Omega(m)$  cycles  $c$  with  $|c| = O(1)$  edges. Suppose w.l.o.g. that  $n$  is a multiple of 3, and consider a tripartite complete graph  $K_{n/3}^3 = (V_1 \cup V_2 \cup V_3, E)$ , where  $V_1, V_2, V_3$  are pairwise disjoint sets of vertices of size  $n/3$  each, and  $E$  is the set of edges thus formed: for each choice of  $x_1 \in V_1, x_2 \in V_2, x_3 \in V_3$ , edges  $(x_1, x_2)$ ,  $(x_2, x_3)$ , and  $(x_1, x_3)$  belong to  $E$ ; and no edges exist that connect any two vertices within  $V_i$ , for  $i = 1, 2, 3$ . Note that each choice of  $x_1 \in V_1, x_2 \in V_2, x_3 \in V_3$  in  $K_{n/3}^3$  gives rise to a distinct cycle  $c \equiv (x_1, x_2), (x_2, x_3), (x_3, x_1)$  of length  $|c| = 3$ , and there are no other (simple) cycles. Thus, there are  $\eta = (n/3)^3$  cycles in the graph  $K_{n/3}^3$ , where  $m = 3(n/3)^2$ . For each graph  $K_{n/3}^3$  in this family, our algorithm requires  $\Theta(\eta + m) = \Theta(\eta) = \Theta(n^3)$  time to list the cycles, thus saving  $O(m + n) = O(n^2)$  time when compared to Johnson's algorithm that takes  $O(n^5)$  in this case.

Note that the analysis of the time complexity of Johnson's algorithm is not pessimistic and cannot match the one of our algorithm for listing cycles. For example, consider the sparse "diamond" graph  $D_n = (V, E)$  in Fig. 1 with  $n = 2k + 3$  vertices in  $V = \{a, b, c, v_1, \dots, v_k, u_1, \dots, u_k\}$ . There are  $m = \Theta(n)$  edges in  $E = \{(a, c), (a, v_i), (v_i, b), (b, u_i), (u_i, c), \text{ for } 1 \leq i \leq k\}$ , and three kinds of (simple) cycles: (1)  $(a, v_i), (v_i, b), (b, u_j), (u_j, c), (c, a)$  for  $1 \leq i, j \leq k$ ; (2)  $(a, v_i), (v_i, b), (b, v_j), (v_j, a)$  for  $1 \leq i < j \leq k$ ; (3)  $(b, u_i), (u_i, c), (c, u_j), (u_j, b)$  for  $1 \leq i < j \leq k$ , totalizing  $\eta = \Theta(n^2)$  cycles. Our algorithm takes  $\Theta(n + k^2) = \Theta(\eta) = \Theta(n^2)$  time to list these cycles. On the other hand, Johnson's algorithm takes  $\Theta(n^3)$  time, and the discovery of the  $\Theta(n^2)$  cycles in (1) costs  $\Theta(k) = \Theta(n)$  time each: the backtracking procedure in Johnson's algorithm starting at  $a$ , and passing through  $v_i, b$  and  $u_j$  for some  $i, j$ , arrives at  $c$ : at that point, it explores all the vertices  $u_l$  ( $l \neq i$ ) even if they do not lead to cycles when coupled with  $a, v_i, b, u_j$ , and  $c$ .

## 2 Overview and Main Ideas

### Preliminaries

Let  $G = (V, E)$  be an undirected connected graph with  $n = |V|$  vertices and  $m = |E|$  edges, without self-loops or parallel edges. For a vertex  $u \in V$ , we denote by  $N(u)$  the neighborhood of  $u$  and by  $d(u) = |N(u)|$  its degree.  $G[V']$  denotes the subgraph *induced* by  $V' \subseteq V$ , and  $G - u$  is the induced subgraph  $G[V \setminus \{u\}]$  for  $u \in V$ . Likewise for edge  $e \in E$ , we adopt the notation  $G - e = (V, E \setminus \{e\})$ .

Paths are simple in  $G$  by definition: we refer to a path  $\pi$  by its natural sequence of vertices or edges. A path  $\pi$  from  $s$  to  $t$ , or *st-path*, is denoted by  $\pi = s \rightsquigarrow t$ . Additionally,  $\mathcal{P}(G)$  is the set of all paths in  $G$  and  $\mathcal{P}_{s,t}(G)$  is the set of all *st*-paths in  $G$ . When  $s = t$  we have cycles, and  $\mathcal{C}(G)$  denotes the set of all cycles in  $G$ . We denote the number of edges in a path  $\pi$  by  $|\pi|$  and in a cycle  $c$  by  $|c|$ . In this paper, we consider the following problems.

**Problem 1 (Listing st-Paths)** *Given a graph  $G = (V, E)$  and two distinct vertices  $s, t \in V$ , output all the paths  $\pi \in \mathcal{P}_{s,t}(G)$ .*

**Problem 2 (Listing Cycles)** *Given a graph  $G = (V, E)$ , output all the cycles  $c \in \mathcal{C}(G)$ .*

Our algorithms assume without loss of generality that the input graph  $G$  is connected, hence  $m \geq n - 1$ , and use the decomposition of  $G$  into biconnected components. Recall that an *articulation point* (or cut-vertex) is a vertex  $u \in V$  such that the number of connected components in  $G$  increases




Figure 1: Diamond graph.




Figure 2: Block tree of  $G$  with bead string  $B_{s,t}$  in gray.

when  $u$  is removed.  $G$  is *biconnected* if it has no articulation points. Otherwise,  $G$  can always be decomposed into a tree of biconnected components, called the *block tree*, where each biconnected component is a maximal biconnected subgraph of  $G$  (see Fig. 2), and two biconnected components are adjacent if and only if they share an articulation point.

## 2.1 Reduction to $st$ -paths

We now show that listing cycles reduces to listing  $st$ -paths while preserving the optimal complexity.

**Lemma 3** *Given an algorithm that solves Problem 1 in optimal  $O(m + \sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$  time, there exists an algorithm that solves Problem 2 in optimal  $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$  time.*

*Proof:* Compute the biconnected components of  $G$  and keep them in a list  $L$ . Each (simple) cycle is contained in one of the biconnected components and therefore we can treat each biconnected component individually as follows. While  $L$  is not empty, extract a biconnected component  $B = (V_B, E_B)$  from  $L$  and repeat the following three steps: (i) compute a DFS traversal of  $B$  and take any back edge  $b = (s, t)$  in  $B$ ; (ii) list all  $st$ -paths in  $B - b$ , i.e. the cycles in  $B$  that include edge  $b$ ; (iii) remove edge  $b$  from  $B$ , compute the new biconnected components thus created by removing edge  $b$ , and append them to  $L$ . When  $L$  becomes empty, all the cycles in  $G$  have been listed.

Creating  $L$  takes  $O(m)$  time. For every  $B \in L$ , steps (i) and (iii) take  $O(|E_B|)$  time. Note that step (ii) always outputs distinct cycles in  $B$  (i.e.  $st$ -paths in  $B - b$ ) in  $O(|E_B| + \sum_{\pi \in \mathcal{P}_{s,t}(B-b)} |\pi|)$  time. However,  $B - b$  is then decomposed into biconnected components whose edges are traversed again. We can pay for the latter cost: for any edge  $e \neq b$  in a biconnected component  $B$ , there is always a cycle in  $B$  that contains both  $b$  and  $e$  (i.e. it is an  $st$ -path in  $B - b$ ), hence  $\sum_{\pi \in \mathcal{P}_{s,t}(B-b)} |\pi|$  dominates the term  $|E_B|$ , i.e.  $\sum_{\pi \in \mathcal{P}_{s,t}(B-b)} |\pi| = \Omega(|E_B|)$ . Therefore steps (i)–(iii) take  $O(\sum_{\pi \in \mathcal{P}_{s,t}(B-b)} |\pi|)$  time. When  $L$  becomes empty, the whole task has taken  $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$  time.  $\square$

## 2.2 Decomposition into biconnected components

We now focus on listing  $st$ -paths (Problem 1). We use the decomposition of  $G$  into a block tree of biconnected components. Given vertices  $s, t$ , define its *bead string*, denoted by  $B_{s,t}$ , as the unique sequence of one or more adjacent biconnected components (the *beads*) in the block tree, such that the first one contains  $s$  and the last one contains  $t$  (see Fig. 2): these biconnected components are connected through articulation points, which must belong to all the paths to be listed.

**Lemma 4** *All the  $st$ -paths in  $\mathcal{P}_{s,t}(G)$  are contained in the induced subgraph  $G[B_{s,t}]$  for the bead string  $B_{s,t}$ . Moreover, all the articulation points in  $G[B_{s,t}]$  are traversed by each of these paths.*

*Proof:* Consider an edge  $e = (u, v)$  in  $G$  such that  $u \in B_{s,t}$  and  $v \notin B_{s,t}$ . Since the biconnected components of a graph form a tree and the bead string  $B_{s,t}$  is a path in this tree, there are no

paths  $v \rightsquigarrow w$  in  $G - e$  for any  $w \in B_{s,t}$  because the biconnected components in  $G$  are maximal and there would be a larger one (a contradiction). Moreover, let  $B_1, B_2, \dots, B_r$  be the biconnected components composing  $B_{s,t}$ , where  $s \in B_1$  and  $t \in B_r$ . If there is only one biconnected component in the path (i.e.  $r = 1$ ), there are no articulation points in  $B_{s,t}$ . Otherwise, all of the  $r - 1$  articulation points in  $B_{s,t}$  are traversed by each path  $\pi \in \mathcal{P}_{s,t}(G)$ : indeed, the articulation point between adjacent biconnected components  $B_i$  and  $B_{i+1}$  is their only vertex in common and there are no edges linking  $B_i$  and  $B_{i+1}$ .  $\square$

We thus restrict the problem of listing the paths in  $\mathcal{P}_{s,t}(G)$  to the induced subgraph  $G[B_{s,t}]$ , conceptually isolating it from the rest of  $G$ . For the sake of description, we will use interchangeably  $B_{s,t}$  and  $G[B_{s,t}]$  in the rest of the paper.

### 2.3 Binary partition scheme

We list the set of  $st$ -paths in  $B_{s,t}$ , denoted by  $\mathcal{P}_{s,t}(B_{s,t})$ , by applying the binary partition method (where  $\mathcal{P}_{s,t}(G) = \mathcal{P}_{s,t}(B_{s,t})$  by Lemma 4): We choose an edge  $e = (s, v)$  incident to  $s$  and then list all the  $st$ -paths that include  $e$  and then all the  $st$ -paths that do not include  $e$ . Since we delete some vertices and some edges during the recursive calls, we proceed as follows.

*Invariant:* At a generic recursive step on vertex  $u$  (initially,  $u := s$ ), let  $\pi_s = s \rightsquigarrow u$  be the path discovered so far (initially,  $\pi_s$  is empty  $\{\}$ ). Let  $B_{u,t}$  be the current bead string (initially,  $B_{u,t} := B_{s,t}$ ). More precisely,  $B_{u,t}$  is defined as follows: (i) remove from  $B_{s,t}$  all the nodes in  $\pi_s$  but  $u$ , and the edges incident to  $u$  and discarded so far; (ii) recompute the block tree on the resulting graph; (iii)  $B_{u,t}$  is the unique bead string that connects  $u$  to  $t$  in the recomputed block tree.

*Base case:* When  $u = t$ , output the  $st$ -path  $\pi_s$  computed so far.

*Recursive rule:* Let  $\mathcal{P}(\pi_s, u, B_{u,t})$  denote the set of  $st$ -paths to be listed by the current recursive call. Then, it is the union of the following two disjoint sets, for an edge  $e = (u, v)$  incident to  $u$ :

- [left branching] the  $st$ -paths in  $\mathcal{P}(\pi_s \cdot e, v, B_{v,t})$  that use  $e$ , where  $B_{v,t}$  is the unique bead string connecting  $v$  to  $t$  in the block tree resulting from the deletion of vertex  $u$  from  $B_{u,t}$ ;
- [right branching] the  $st$ -paths in  $\mathcal{P}(\pi_s, u, B'_{u,t})$  that do *not* use  $e$ , where  $B'_{u,t}$  is the unique bead string connecting  $u$  to  $t$  in the block tree resulting from the deletion of edge  $e$  from  $B_{u,t}$ .

Hence,  $\mathcal{P}_{s,t}(B_{s,t})$  (and so  $\mathcal{P}_{s,t}(G)$ ) can be computed by invoking  $\mathcal{P}(\{\}, s, B_{s,t})$ . The correctness and completeness of the above approach is discussed in Section 2.4.

At this point, it should be clear why we introduce the notion of bead strings in the binary partition. The existence of the partial path  $\pi_s$  and the bead string  $B_{u,t}$  guarantees that there surely exists at least one  $st$ -path. But there are two sides of the coin when using  $B_{u,t}$ .

(1) One advantage is that we can avoid useless recursive calls: If vertex  $u$  has only one incident edge  $e$ , we just perform the left branching; otherwise, we can safely perform both the left and right branching since the *first* bead in  $B_{u,t}$  is always a biconnected component by definition (and so there exist both an  $st$ -path that traverses  $e$  and one that does not traverse  $e$ ).

(2) The other side of the coin is that we have to maintain the bead string  $B_{u,t}$  as  $B_{v,t}$  in the left branching and as  $B'_{u,t}$  in the right branching by Lemma 4. Note that these bead strings are surely non-empty since  $B_{u,t}$  is non-empty by induction (we only perform either left or left/right branching when there are solutions by point (1)).

To efficiently address point (2), we need to introduce the notion of certificate as described next.




Figure 3: Example certificate  $C$

---

**Algorithm 1**  $\text{list\_paths}_{s,t}(\pi_s, u, C)$

---

```

1: if  $u = t$  then
2:   output( $\pi_s$ )
3:   return
4: end if
5:  $e = (u, v) := \text{choose}(C, u)$ 
6: if  $e$  is back edge then
7:    $I := \text{right\_update}(C, e)$ 
8:    $\text{list\_paths}_{s,t}(\pi_s, u, C)$ 
9:    $\text{restore}(C, I)$ 
10: end if
11:  $I := \text{left\_update}(C, e)$ 
12:  $\text{list\_paths}_{s,t}(\pi_s \cdot (u, v), v, C)$ 
13:  $\text{restore}(C, I)$ 

```

---

## 2.4 Introducing the certificate

Given the bead string  $B_{u,t}$ , we call the *head* of  $B_{u,t}$ , denoted by  $H_u$ , the first biconnected component in  $B_{u,t}$ , where  $u \in H_u$ . Consider a DFS tree of  $B_{u,t}$  rooted at  $u$  that changes along with  $B_{u,t}$ , and classify the edges in  $B_{u,t}$  as tree edges or back edges (no cross edges since the graph is undirected).

To maintain  $B_{u,t}$  (and so  $H_u$ ) during the recursive calls, we introduce a *certificate*  $C$  (see Fig. 3 for an example): It is a suitable data structure that uses the above classification of the edges in  $B_{u,t}$ , and supports the following operations, required by the binary partition scheme.

- **choose**( $C, u$ ): returns an edge  $e = (u, v)$  with  $v \in H_u$  such that  $\pi_s \cdot (u, v) \cdot u \rightsquigarrow t$  is an  $st$ -path such that  $u \rightsquigarrow t$  is inside  $B_{u,t}$ . Note that  $e$  always exists since  $H_u$  is biconnected. Also, the chosen  $v$  is the last one in DFS order among the neighbors of  $u$ : in this way, the (only) tree edge  $e$  is returned when there are no back edges leaving from  $u$ .<sup>1</sup>
- **left\_update**( $C, e$ ): for the given  $e = (u, v)$ , it obtains  $B_{v,t}$  from  $B_{u,t}$  as discussed in Section 2.3. This implies updating also  $H_u$ ,  $C$ , and the block tree, since the recursion continues on  $v$ . It returns bookkeeping information  $I$  for what is updated, so that it is possible to revert to  $B_{u,t}$ ,  $H_u$ ,  $C$ , and the block tree, to their status before this operation.
- **right\_update**( $C, e$ ): for the given  $e = (u, v)$ , it obtains  $B'_{u,t}$  from  $B_{u,t}$  as discussed in Section 2.3, which implies updating also  $H_u$ ,  $C$ , and the block tree. It returns bookkeeping information  $I$  as in the case of **left\_update**( $C, e$ ).
- **restore**( $C, I$ ): reverts the bead string to  $B_{u,t}$ , the head  $H_u$ , the certificate  $C$ , and the block tree, to their status before operation  $I := \text{left\_update}(C, e)$  or  $I := \text{right\_update}(C, e)$  was issued (in the same recursive call).

Note that a notion of certificate in listing problems has been introduced in [4], but it cannot be directly applied to our case due to the different nature of the problems and our use of more complex structures such as biconnected components. Using our certificate and its operations, we can now formalize the binary partition and its recursive calls  $\mathcal{P}(\pi_s, u, B_{u,t})$  described in Section 2.3 as Algorithm 1, where  $B_{u,t}$  is replaced by its certificate  $C$ .

---

<sup>1</sup>As it will be clear in Sections 3 and 4, this order facilitates the analysis and the implementation of the certificate.

The base case ( $u = t$ ) corresponds to lines 1–4 of Algorithm 1. During recursion, the left branching corresponds to lines 5 and 11–13, while the right branching to lines 5–10. Note that we perform only the left branching when there is only one incident edge in  $u$ , which is a tree edge by definition of **choose**. Also, lines 9 and 13 are needed to restore the parameters to their values when returning from the recursive calls. The proof of the following lemma is in the Appendix.

**Lemma 5** *Algorithm 1 correctly lists all the  $st$ -paths in  $\mathcal{P}_{s,t}(G)$ .*

A natural question is what is the complexity: we should account for the cost of maintaining  $C$  and for the cost of the recursive calls of Algorithm 1. Since we cannot always maintain the certificate in  $O(1)$  time, the ideal situation for attaining an optimal cost is taking  $O(\mu)$  time if at least  $\mu$   $st$ -paths are listed in the current call (and its nested calls). An obstacle to this ideal situation is that we cannot estimate  $\mu$  efficiently and cannot design Algorithm 1 so that it takes  $O(\mu)$  adaptively. We circumvent this by using a different cost scheme in Section 2.5 that is based on the recursion tree induced by Algorithm 1. Section 4 and Appendix B are devoted to the efficient implementation of the above certificate operations according to the cost scheme that we discuss next.

## 2.5 Recursion tree and cost amortization

We now define how to distribute the costs among the several recursive calls of Algorithm 1 so that optimality is achieved. Consider a generic execution on the bead string  $B_{u,t}$ . We can trace this execution by using a binary recursion tree  $R$ . The nodes of  $R$  are labeled by the arguments on which Algorithm 1 is run: specifically, we denote a node in  $R$  by the triple  $x = \langle \pi_s, u, C \rangle$  iff it represents the call with arguments  $\pi_s$ ,  $u$ , and  $C$ .<sup>2</sup> The left branching is represented by the left child, and the right branching (if any) by the right child of the current node.

**Lemma 6** *The binary recursion tree  $R$  for  $B_{u,t}$  has the following properties:*

1. *There is a one-to-one correspondence between the paths in  $\mathcal{P}_{s,t}(B_{u,t})$  and the leaves in the recursion tree rooted at node  $\langle \pi_s, u, C \rangle$ .*
2. *Consider any leaf and its corresponding  $st$ -path  $\pi$ : there are  $|\pi|$  left branches in the corresponding root-to-leaf trace.*
3. *Consider the instruction  $e := \text{choose}(C, u)$  in Algorithm 1: unary (i.e. single-child) nodes correspond to left branches ( $e$  is a tree edge) while binary nodes correspond to left and right branches ( $e$  is a back edge).*
4. *The number of binary nodes is  $|\mathcal{P}_{s,t}(B_{u,t})| - 1$ .*

We define a *spine* of  $R$  to be a subset of  $R$ 's nodes linked as follows: the first node is a node  $x$  that is either the left child of its parent or the root of  $R$ , and the other nodes are those reachable from  $x$  by right branching in  $R$ . Let  $x = \langle \pi_s, u, C \rangle$  be the first node in a spine  $S$ . The nodes in  $S$  correspond to the edges that are incident to vertex  $u$  in  $B_{u,t}$ : hence their number equals the degree  $d(u)$  of  $u$  in  $B_{u,t}$ , and the deepest (last) node in  $S$  is always a tree edge in  $B_{u,t}$  while the others are back edges (in reverse DFS order). Summing up,  $R$  can be seen as composed by spines, unary nodes, and leaves, where each spine has a unary node as deepest node. This gives a global picture of  $R$  that we now exploit for the analysis.

---

<sup>2</sup>For clarity, we use “nodes” when referring to  $R$  and “vertices” when referring to  $B_{u,t}$ .

We define the *compact head*, denoted by  $H_X = (V_X, E_X)$ , as the (multi)graph obtained by compacting the maximal chains of degree-2 vertices, except  $u$ ,  $t$ , and the vertices that are the leaves of its DFS tree rooted at  $u$ .

The rationale behind the above definition is that the costs defined in terms of  $H_X$  amortize well, as the size of  $H_X$  and the number of  $st$ -paths in the subtree of  $R$  rooted at node  $x = \langle \pi_s, u, C \rangle$  are intimately related (see Lemma 10 in Section 3) while this is not necessarily true for  $H_u$ .

We now define the following abstract cost for spines, unary nodes, and leaves of  $R$ , for a sufficiently large constant  $c_0 > 0$ , that Algorithm 1 must fulfill:

$$T(r) = \begin{cases} c_0 & \text{if } r \text{ is a unary node} \\ c_0|\pi| & \text{if } r \text{ is a leaf corresponding to path } \pi \\ c_0(|V_X| + |E_X|) & \text{if } r \text{ is a spine with compact head } H_X \end{cases} \quad (1)$$

**Lemma 7** *The sum of the costs in the nodes of the recursion tree  $\sum_{r \in R} T(r) = O(\sum_{\pi \in \mathcal{P}_{s,t}(B_{u,t})} |\pi|)$ .*

Section 3 contains the proof of Lemma 7 and related properties. Setting  $u := s$ , we obtain that the cost in Lemma 7 is optimal, by Lemma 4.

**Theorem 8** *Algorithm 1 solves problem Problem 1 in optimal  $O(m + \sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$  time.*

By Lemma 3, we obtain an optimal result for listing cycles.

**Theorem 9** *Problem 2 can be optimally solved in  $O(m + \sum_{c \in \mathcal{C}(G)} |c|)$  time.*

### 3 Amortization strategy

We devote this section to prove Lemma 7. Let us split the sum in Eq. (1) in three parts, and bound each part individually, as

$$\sum_{r \in R} T(r) \leq \sum_{r: \text{unary}} T(r) + \sum_{r: \text{leaf}} T(r) + \sum_{r: \text{spine}} T(r). \quad (2)$$

We have that  $\sum_{r: \text{unary}} T(r) = O(\sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$ , since there are  $|\mathcal{P}_{s,t}(G)|$  leaves, and the root-to-leaf trace leading to the leaf for  $\pi$  contains at most  $|\pi|$  unary nodes by Lemma 6, where each unary node has cost  $O(1)$  by Eq. (1).

Also,  $\sum_{r: \text{leaf}} T(r) = O(\sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$ , since the leaf  $r$  for  $\pi$  has cost  $O(|\pi|)$  by Eq. (1).

It remains to bound  $\sum_{r: \text{spine}} T(r)$ . By Eq. (1), we can rewrite this cost as  $\sum_{H_X} c_0(|V_X| + |E_X|)$ , where the sum ranges over the compacted heads  $H_X$  associated with the spines  $r$ . We use the following lemma to provide a lower bound on the number of  $st$ -paths descending from  $r$ .

**Lemma 10** *Given a spine  $r$ , and its bead string  $B_{u,t}$  with head  $H_u$ , there are at least  $|E_X| - |V_X| + 1$   $st$ -paths in  $G$  that have prefix  $\pi_s = s \rightsquigarrow u$  and suffix  $u \rightsquigarrow t$  internal to  $B_{u,t}$ , where the compacted head is  $H_X = (V_X, E_X)$ .*

*Proof:*  $H_X$  is biconnected. In any biconnected graph  $B = (V_B, E_B)$  there are at least  $|E_B| - |V_B| + 1$   $xy$ -paths for any  $x, y \in V_B$ . Find a ear decomposition [3] of  $B$  and consider the process of forming  $B$  by adding ears one at the time, starting from a single cycle including  $x$  and  $y$ . Initially  $|V_B| = |E_B|$  and there are 2  $xy$ -paths. Each new ear forms a path connecting two vertices that are part of a  $xy$ -path, increasing the number of paths by at least 1. If the ear has  $k$  edges, its addition increases  $V$  by  $k - 1$ ,  $E$  by  $k$ , and the number of  $xy$ -paths by at least 1. The result follows by induction.  $\square$



The implication of Lemma 10 is that there are at least  $|E_X| - |V_X| + 1$  leaves descending from the given spine  $r$ . Hence, we can charge to each of them a cost of  $\frac{c_0(|V_X|+|E_X|)}{|E_X|-|V_X|+1}$ . Lemma 11 allows us to prove that the latter cost is  $O(1)$  when  $H_u$  is different from a single edge or a cycle. (If  $H_u$  is a single edge or a cycle,  $H_X$  is a single or double edge, and the cost is trivially a constant.)

**Lemma 11** *For a compacted head  $H_X = (V_X, E_X)$ , its density is  $\frac{|E_X|}{|V_X|} \geq \frac{11}{10}$ .*

Specifically, let  $\alpha = \frac{11}{10}$  and write  $\alpha = 1 + 2/\beta$  for a constant  $\beta$ : we have that  $|E_X| + |V_X| = (|E_X| - |V_X|) + 2|V_X| \leq (|E_X| - |V_X|) + \beta(|E_X| - |V_X|) = \frac{\alpha+1}{\alpha-1}(|E_X| - |V_X|)$ . Thus, we can charge each leaf with a cost of  $\frac{c_0(|V_X|+|E_X|)}{|E_X|-|V_X|+1} \leq c_0 \frac{\alpha+1}{\alpha-1} = O(1)$ . This motivates the definition of  $H_X$ , since Lemma 11 does not necessarily hold for the head  $H_u$  (due to the unary nodes in its DFS tree).

One last step to bound  $\sum_{H_X} c_0(|V_X| + |E_X|)$ : as noted before, a root-to-leaf trace for the string storing  $\pi$  has  $|\pi|$  left branches by Lemma 6, and as many spines, each spine charging  $c_0 \frac{\alpha+1}{\alpha-1} = O(1)$  to the leaf at hand. This means that each of the  $|\mathcal{P}_{s,t}(G)|$  leaves is charged for a cost of  $O(|\pi|)$ , thus bounding the sum as  $\sum_{r \text{ spine}} T(r) = \sum_{H_X} c_0(|V_X| + |E_X|) = O(\sum_{\pi \in \mathcal{P}_{s,t}(G)} |\pi|)$ . This completes the proof of Lemma 7. As a corollary, we obtain the following result.

**Lemma 12** *The recursion tree  $R$  with cost as in Eq. (1) induces an  $O(|\pi|)$  amortized cost for each st-path  $\pi$ .*

## 4 Certificate implementation and maintenance

The certificate  $C$  associated with a node  $\langle \pi_s, u, C \rangle$  in the recursion tree is a compacted and augmented DFS tree of bead string  $B_{u,t}$ , rooted at vertex  $u$ . The DFS tree changes over time along with  $B_{u,t}$ , and is maintained in such a way that  $t$  is in the leftmost path of the tree. We compact the DFS tree by contracting the vertices that have degree 2, except  $u$ ,  $t$ , and the leaves (the latter surely have incident back edges). Maintaining this compacted representation is not a difficult data-structure problem. From now on we can assume w.l.o.g. that  $C$  is an augmented DFS tree rooted at  $u$  where internal nodes of the DFS tree have degree  $\geq 3$ , and each vertex  $v$  has associated:

1. A doubly-linked list  $lb(v)$  of back edges linking  $v$  to its descendants  $w$  sorted by DFS order.
2. A doubly-linked list  $ab(v)$  of back edges linking  $v$  to its ancestors  $w$  sorted by DFS order.
3. An integer  $\gamma(v)$ , such that if  $v$  is an ancestor of  $w$  then  $\gamma(v) < \gamma(w)$ .
4. The smallest  $\gamma(w)$  over all neighbors  $w$  of  $v$  in  $C$ , excluding the parent, denoted by  $lowpoint(v)$ .

Given two vertices  $v, w \in C$  such that  $v$  is the parent of  $w$ , we can efficiently test if any of the children of  $w$  is in the same biconnected component of  $v$ , i.e.  $lowpoint(w) \leq \gamma(v)$ . (Note that we adopt a variant of  $lowpoint$  using  $\gamma(v)$  in place of  $depth(v)$ : it has the same effect whereas using  $\gamma(v)$  is preferable since it is easier to dynamically maintain.)

**Lemma 13** *The certificate associated with the root of the recursion can be computed in  $O(m)$  time.*

*Proof:* In order to set  $t$  to be in the leftmost path, we perform a DFS traversal of graph  $G$  starting from  $s$  and stop when we reach vertex  $t$ . We then compute the DFS tree, traversing the path  $s \rightsquigarrow t$  first. When visiting vertex  $v$ , we set  $\gamma(v)$  to depth of  $v$  in the DFS. Before going up on the traversal, we compute the lowpoints using the lowpoints of the children. Let  $z$  be the parent of  $v$ . If  $lowpoint(v) > lowpoint(z)$  and  $w$  is not in the leftmost path in the DFS, we cut the subtree of

$v$  as it does not belong to  $B_{s,t}$ . When finding a back edge  $e = (v, w)$ , if  $w$  is a descendant of  $v$  we append  $e$  to both  $lb(v)$  and  $ab(w)$ ; else we append  $e$  to both  $ab(v)$  and  $lb(w)$ . This maintains the DFS order in the back edge lists. This procedure takes at most two DFS traversals in  $O(m)$  time. This DFS tree can be compacted in the same time bound.  $\square$

**Lemma 14** *Operation  $\text{choose}(C, u)$  can be implemented in  $O(1)$  time.*

*Proof:* If the list  $lb(v)$  is empty, return the tree edge  $e = (u, v)$  linking  $u$  to its only child  $v$  (there are no other children). Else, return the last edge in  $lb(v)$ .  $\square$

We analyze the cost of updating and restoring the certificate  $C$ . We can reuse parts of  $C$ , namely, those corresponding to the vertices that are not in the compacted head  $H_X = (V_X, E_X)$  as defined in Section 2.5. We prove that, given a unary node  $u$  and its tree edge  $e = (u, v)$ , the subtree of  $v$  in  $C$  can be easily made a certificate for the left branch of the recursion.

**Lemma 15** *On a unary node,  $\text{left\_update}(C, e)$  takes  $O(1)$  time.*

*Proof:* Take edge  $e = (u, v)$ . Remove edge  $e$  and set  $v$  as the root of the certificate. Since  $e$  is the only edge incident in  $v$ , the subtree  $v$  is still a DFS tree. Cut the list of children of  $v$  keeping only the first child. (The other children are no longer in the bead string and become part of  $I$ .) There is no need to update  $\gamma(v)$ .  $\square$

We now devote the rest of this section to show how to efficiently maintain  $C$  on a spine. Consider removing a back edge  $e$  from  $u$ : the compacted head  $H_X = (V_X, E_X)$  of the bead string can be divided into smaller biconnected components. Many of those can be excluded from the certificate (i.e. they are no longer in the new bead string, and so they are bookkept in  $I$ ) and additionally we have to update the lowpoints that change. We prove that this operation can be performed in  $O(|V_X|)$  total time on a spine of the recursion tree.

**Lemma 16** *The total cost of all the operations  $\text{right\_update}(C, e)$  in a spine is  $O(|V_X|)$  time.*

*Proof:* In the right branches along a spine, we remove all back edges in  $lb(u)$ . This is done by starting from the last edge in  $lb(u)$ , i.e. proceeding in reverse DFS order. For back edge  $b_i = (z_i, u)$ , we traverse the vertices in the path from  $z_i$  towards the root  $u$ , as these are the only lowpoints that can change. While moving upwards on the tree, on each vertex  $w$ , we update  $\text{lowpoint}(w)$ . This is done by taking the endpoint  $y$  of the first edge in  $ab(w)$  (the back edge that goes the topmost in the tree) and choosing the minimum between  $\gamma(y)$  and the lowpoint of each child of  $w$ . We stop when the updated  $\text{lowpoint}(w) = \gamma(u)$  since it implies that the lowpoint of the vertex can not be further reduced. Note that we stop before  $u$ , except when removing the last back edge in  $lb(u)$ .

To prune the branches of the DFS tree that are no longer in  $B_{u,t}$ , consider again each vertex  $w$  in the path from  $z_i$  towards the root  $u$  and its parent  $y$ . It is possible to test if  $w$  is an articulation point by checking if the updated  $\text{lowpoint}(w) > \gamma(y)$ . If that is the case and  $w$  is not in the leftmost path of the DFS, it implies that  $w \notin B_{u,t}$ , and therefore we cut the subtree of  $w$  and keep it in  $I$  to restore later. We use the same halting criterion as in the previous paragraph.

The cost of removing all back edges in the spine is  $O(|V_X|)$ : there are  $O(|V_X|)$  tree edges and, in the paths from  $z_i$  to  $u$ , we do not traverse the same tree edge twice since the process described stops at the first common ancestor of endpoints of back edges  $b_i$ . Additionally, we take  $O(1)$  time to cut a subtree of an articulation point in the DFS tree.  $\square$

To compute  $\text{left\_update}(C, e)$  in the binary nodes of a spine, we use the fact that in every left branching from that spine, the graph is the same (in a spine we only remove edges incident to  $u$  and on a left branch from the spine we remove the node  $u$ ) and therefore its block tree is also the same. However, the certificates on these nodes are not the same, as they are rooted at different vertices. By using the DFS order of the edges, we are able to traverse each edge in  $H_X$  only a constant number of times in the spine.

**Lemma 17** *The total cost of all operations  $\text{left\_update}(C, e)$  in a spine is amortized  $O(|E_X|)$ .*

*Proof:* Let  $t'$  be the last vertex in the path  $u \rightsquigarrow t$  s.t.  $t' \in V_X$ . Since  $t'$  is an articulation point, the subtree of the DFS tree rooted in  $t'$  is maintained in the case of removal of vertex  $u$ . Therefore the only modifications of the DFS tree occur in the compacted head  $H_X$  of  $B_{u,t}$ . Let us compute the certificate  $C_i$ : this is the certificate of the left branch of the  $i$ th node of the spine where we augment the path with the back edge  $b_i = (z_i, u)$  of  $lb(u)$  in the order defined by  $\text{choose}(C, u)$ .

For the case of  $C_1$ , we remove  $u$  and rebuild the certificate starting from  $z_1$  (the last edge in  $lb(u)$ ) using the algorithm from Lemma 13 restricted to  $H_X$  and using  $t'$  as target and  $\gamma(t')$  as a baseline to  $\gamma$  (instead of the depth). This takes  $O(|E_X|)$ .

For the general case of  $C_i$  with  $i > 1$  we also rebuild (part) of the certificate starting from  $z_i$  using the procedure from Lemma 13 but we use information gathered in  $C_{i-1}$  to avoid exploring useless branches of the DFS tree. The key point is that, when we reach the first bead in common to both  $B_{z_i,t}$  and  $B_{z_{i-1},t}$ , we only explore edges internal to this bead. If an edge  $e$  leaving the bead leads to  $t$ , we can reuse a subtree of  $C_{i-1}$ . If  $e$  does not lead to  $t$ , then it has already been explored (and cut) in  $C_{i-1}$  and there is no need to explore it again since it will be discarded. Given the order we take  $b_i$ , each bead is not added more than once, and the total cost over the spine is  $O(|E_X|)$ .

Nevertheless, the internal edges  $E'_X$  of the first bead in common between  $B_{z_i,t}$  and  $B_{z_{i-1},t}$  can be explored several times during this procedure.<sup>3</sup> We can charge the cost  $O(|E'_X|)$  of exploring those edges to another node in the recursion tree, since this common bead is the head of at least one certificate in the recursion subtree of the left child of the  $i$ th node of the spine. Specifically, we charge the first node in the *leftmost* path of the  $i$ th node of the spine that has exactly the edges  $E'_X$  as head of its bead string: (i) if  $|E'_X| \leq 1$  it corresponds to a unary node or a leaf in the recursion tree and therefore we can charge it with  $O(1)$  cost; (ii) otherwise it corresponds to a first node of a spine and therefore we can also charge it with  $O(|E'_X|)$ . We use this charging scheme when  $i \neq 1$  and the cost is always charged in the leftmost recursion path of  $i$ th node of the spine. Consequently, we never charge a node in the recursion tree more than once.  $\square$

**Lemma 18** *On each node of the recursion tree,  $\text{restore}(C, I)$  takes time proportional to the size of the modifications kept in  $I$ .*

From Lemmas 14 and 16–18, it follows that on a spine of the recursion tree we have the costs:  $\text{choose}(u)$  on each node which is bounded by  $O(|V_X|)$  as there are at most  $|V_X|$  back edges in  $u$ ;  $\text{right\_update}(C, e)$ ,  $\text{restore}(C, I)$  take  $O(|V_X|)$  time;  $\text{left\_update}(C, e)$  and  $\text{restore}(C, I)$  are charged  $O(|V_X| + |E_X|)$  time. We thus have the following result, completing the proof of Theorem 8.

**Lemma 19** *Algorithm 1 can be implemented with a cost fulfilling Eq. (1), thus it takes total  $O(m + \sum_{r \in R} T(r)) = O(m + \sum_{\pi \in \mathcal{P}_{s,t}(B_{u,t})} |\pi|)$  time.*

---

<sup>3</sup>Consider the case where  $z_i, \dots, z_j$  are all in the same bead after the removal of  $u$ . The bead strings are the same, but the roots  $z_i, \dots, z_j$  are different, so we have to compute the corresponding DFS of the first component  $|j - i|$  times.

## References

- [1] G.J. Bezem and J. van Leeuwen. Enumeration in graphs. Technical Report RUU-CS-87-07, Utrecht University, 1987.
- [2] E. Birmelé, P. Crescenzi, R. Ferreira, R. Grossi, V. Lacroix, A. Marino, N. Pisanti, G. Sacomoto, and M.-F. Sagot. Efficient bubble enumeration in directed graphs. In *Proc. String Processing and Information Retrieval, 19th International Symposium, SPIRE 2012*, Lecture Notes in Computer Science. Springer, 2012 (to appear).
- [3] Reinhard Diestel. *Graph Theory (Graduate Texts in Mathematics)*. Springer, 2005.
- [4] Rui A. Ferreira, Roberto Grossi, and Romeo Rizzi. Output-sensitive listing of bounded-size trees in undirected graphs. In *ESA*, pages 275–286, 2011.
- [5] T. R. Halford and K. M. Chugg. Enumerating and counting cycles in bipartite graphs. In *IEEE Communication Theory Workshop*, 2004.
- [6] Tamás Horváth, Thomas Gärtner, and Stefan Wrobel. Cyclic pattern kernels for predictive graph mining. In *Proc. of 10th ACM SIGKDD*, pages 158–167, 2004.
- [7] Donald B. Johnson. Finding all the elementary circuits of a directed graph. *SIAM J. Comput.*, 4(1):77–84, 1975.
- [8] S. Klamt and et al. A methodology for the structural and functional analysis of signaling and regulatory networks. *BMC Bioinformatics*, 7:56, 2006.
- [9] S. Klamt and A. von Kamp. Computing paths and cycles in biological interaction graphs. *BMC Bioinformatics*, 10:181, 2009.
- [10] Hongbo Liu and Jiaxin Wang. A new way to enumerate cycles in graph. In *AICT and ICIW*, pages 57–59, 2006.
- [11] Prabhakar Mateti and Narsingh Deo. On algorithms for enumerating all circuits of a graph. *SIAM J. Comput.*, 5(1):90–99, 1976.
- [12] J. Poinstein. Self-avoiding paths and the adjacency matrix of a graph. *SIAM Journal on Applied Mathematics*, 14:600–609, 1966.
- [13] R C Read and Robert E Tarjan. Bounds on backtrack algorithms for listing cycles, paths, and spanning trees. *Networks*, 5(3):237–252, 1975.
- [14] K. Sankar and A.V. Sarad. A time and memory efficient way to enumerate cycles in a graph. In *Intelligent and Advanced Systems*, pages 498–500, 2007.
- [15] R. Schott and George Stacey Staples. Complexity of counting cycles using Zeons. *Computers and Mathematics with Applications*, 62:1828–1837, 2011.
- [16] E.H. Sussenguth. A graph-theoretical algorithm for matching chemical structures. *J. Chem. Doc.*, 5:36–43, 1965.
- [17] Maciej M. Syslo. An efficient cycle vector space algorithm for listing all cycles of a planar graph. *SIAM J. Comput.*, 10(4):797–808, 1981.
- [18] Jayme L. Szwarcfiter and Peter E. Lauer. A search strategy for the elementary cycles of a directed graph. *BIT Numerical Mathematics*, 16, 1976.
- [19] Robert Endre Tarjan. Enumeration of the elementary circuits of a directed graph. *SIAM J. Comput.*, 2(3):211–216, 1973.
- [20] James C. Tiernan. An efficient search algorithm to find the elementary circuits of a graph. *Communications ACM*, 13:722–726, 1970.
- [21] John T. Welch, Jr. A mechanical analysis of the cyclic structure of undirected linear graphs. *J. ACM*, 13:205–210, 1966.
- [22] Marcel Wild. Generating all cycles, chordless cycles, and hamiltonian cycles with the principle of exclusion. *J. of Discrete Algorithms*, 6:93–102, 2008.
- [23] S.S. Yau. Generation of all hamiltonian circuits, paths, and centers of a graph, and related problems. *IEEE Transactions on Circuit Theory*, 14:79–81, 1967.

## A Omitted Proofs

### A.1 Lemma 5

*Proof:* For a given vertex  $u$  the function `choose`( $C, u$ ) returns an edge  $e$  incident to  $u$ . We maintain the invariant that  $\pi_s$  is a path  $s \rightsquigarrow u$ , since at the point of the recursive call in line 12: (i) is connected as we append edge  $(u, v)$  to  $\pi_s$  and; (ii) it is simple as vertex  $u$  is removed from the graph  $G$  in the call to `left_update`( $C, e$ ) in line 11. In the case of recursive call in line 8 the invariant is trivially maintained as  $\pi_s$  does not change. The algorithm only outputs  $st$ -paths since  $\pi_s$  is a  $s \rightsquigarrow u$  path and  $u = t$  when the algorithm outputs, in line 2.

All the paths with prefix  $\pi_s$  that do not use  $e$  are listed by the recursive call in line 8. This is done by removing  $e$  from the graph in line 7 and thus no path can include  $e$ . The paths that use  $e$  are listed in line 12 since in the recursive call  $e$  is added to  $\pi_s$ . Given that the tree edge incident to  $u$  is the last one to be returned by `choose`( $C, u$ ), there is no path that does not use this edge, therefore it is not necessary to call line 8 for this edge.  $\square$

### A.2 Lemma 6

*Proof:* We proceed in order as follows.

1. We only output a solution in a leaf and we only do recursive calls that lead us to a solution. Moreover every node partitions the set of solutions in the ones that use an edge and the ones that do not use it. This guarantees that the leaves in the left subtree of the node corresponding to the recursive call and the leaves in the right subtree do not intersect. This implies that different leaves correspond to different paths from  $s$  to  $t$ , and that for each path there is a corresponding leaf.
2. Each left branch corresponds to the inclusion of an edge in the path  $\pi$ .
3. Since we are in a biconnected component, there is always a left branch. There can be no unary node as a right branch: indeed for any edge of  $B_{u,t}$  there exists always a path from  $s$  to  $t$  passing through that edge. Since the tree edge is always the last one to be chosen, unary nodes cannot correspond to back edges and binary nodes are always back edges.
4. From point 1 and from the fact that the recursion tree is a binary tree. (In any binary tree, the number of binary nodes is equal to the number of leaves minus 1.)

$\square$

### A.3 Lemma 11

*Proof:* Consider the following partition  $V_X = \{r\} \cup V_2 \cup V_3$  where:  $r$  is the root;  $V_2$  is the set of vertices with degree 2 and;  $V_3$ , the vertices with degree  $\geq 3$ . Since  $H_X$  is compacted DFS tree of a biconnected graph, we have that  $V_2$  is a *subset* of the leaves and  $V_3$  contains the set of internal nodes (except  $r$ ). There are no vertices with degree 1 and  $d(r) \geq 2$ . Let  $x = \sum_{v \in V_3} d(v)$  and  $y = \sum_{v \in V_2} d(v)$ . We can write the density as a function of  $x$  and  $y$ :

$$\frac{|E_X|}{|V_X|} = \frac{x + y + d(r)}{2(|V_3| + |V_2| + 1)}$$

Note that  $|V_3| \leq \frac{x}{3}$  as the vertices in  $V_3$  have at least degree 3,  $|V_2| = \frac{y}{2}$  as vertices in  $V_2$  have degree exactly 2. Since  $d(r) \geq 2$ , we derive the following bound:

$$\frac{|E_X|}{|V_X|} \geq \frac{x + y + 2}{\frac{2}{3}x + y + 2}$$

Consider any graph with  $|V_X| > 3$  and its DFS tree rooted at  $r$ . Note that: (i) there are no tree edges between any two leaves, (ii) every node in  $V_2$  is a leaf and (iii) no leaf is a child of  $r$ . Therefore, every tree edge incident in a vertex of  $V_2$  is also incident in a vertex of  $V_3$ . Since exactly half the incident edges to  $V_2$  are tree edges (the other half are back edges) we get that  $y \leq 2x$ .

With  $|V_X| \geq 3$  there exists at least one internal node in the DFS tree and therefore  $x \geq 3$ .

$$\begin{aligned} & \text{minimize} && \frac{x + y + 2}{\frac{2}{3}x + y + 2} \\ & \text{subject to} && 0 \leq y \leq 2x, \\ & && x \geq 3. \end{aligned}$$

Since for any  $x$  the function is minimized by the maximum  $y$  s.t.  $y \leq 2x$  and for any  $y$  by the minimum  $x$ , we get:

$$\frac{|E_X|}{|V_X|} \geq \frac{9x + 6}{8x + 6} \geq \frac{11}{10}$$

□

#### A.4 Lemma 18

*Proof:* We use standard data structures (i.e. linked lists) for the representation of certificate  $C$ . Persistent versions of these data structures exist that maintain a stack of modifications applied to them and that can restore its contents to their previous states. Given the modifications in  $I$ , these data structures take  $O(|I|)$  time to restore the previous version of  $C$ .

Let us consider the case of performing `left_update`( $C, e$ ). We cut at most  $O(|V_X|)$  edges from  $C$ . Note that, although we conceptually remove whole branches of the DFS tree, we only remove edges that attach those branches to the DFS tree. The other vertices and edges are left in the certificate but, as they no longer remain attached to  $B_{u,t}$ , they will never be reached or explored. In the case of `right_update`( $C, e$ ), we have a similar situation, with at most  $O(|E_X|)$  edges being modified along the spine of the recursion tree. □




Figure 4: Example bead string and spine of the recursion tree

## B Extended analysis of operations in a spine of the recursion tree

In this appendix, we present all details and illustrate with figures the operations `right_update`( $C, e$ ) and `left_update`( $C, e$ ) that are performed along a spine of the recursion tree. In order to better detail the procedures in Lemma 16 and Lemma 17, we divide them in smaller parts. Fig. 4 shows (a) an example of a bead string  $B_{u,t}$  at the first node of the spine and (b) the nodes of the spine. This spine contains four binary nodes corresponding to the back edges in  $lb(u)$  and an unary node corresponding to the tree edge  $(u, v)$ . Note that edges are taken in reverse DFS order as defined in operation `choose`( $C, u$ ).

As a consequence of Lemma 4, the impact of operations `right_update`( $C, e$ ) and `left_update`( $C, e$ ) in the certificate is restricted to the biconnected component of  $u$ . Thus we mainly focus on maintaining the compacted head  $H_X = (V_X, E_X)$  of the bead string  $B_{u,t}$ .

### B.1 Operation `right_update`( $C, e$ ) in a spine of the recursion tree

Let us now prove and illustrate the following lemma:

**Lemma 20** (Lemma 16 restated) *In a spine of the recursion tree, operations `right_update`( $C, e$ ) can be implemented in  $O(|V_X|)$  total time.*

In the right branches along a spine, we remove all back edges in  $lb(u)$ . This is done by starting from the last edge in  $lb(u)$ , i.e. proceeding in reverse DFS order. In the example from Fig. 4, we remove the back edges  $(z_1, u)$ ,  $(z_2, u)$ ,  $(z_3, u)$  and  $(z_4, u)$ . To update the certificate corresponding to  $B_{u,t}$ , we have to (i) update the lowpoints in each vertex of  $H_X$  (ii) prune vertices that are no

longer in  $B_{u,t}$  after removing a back edge. Note that, for a vertex  $w$  in the tree, there is no need to update  $\gamma(w)$ .

Let us consider the update of lowpoints in the DFS tree. For a back edge  $b_i = (z_i, u)$ , we traverse the vertices in the path from  $z_i$  towards the root  $u$ . By definition of lowpoint, these are the only lowpoints that can change. Suppose that we remove back edge  $(z_4, u)$  in the example from Fig. 4, only the lowpoints of the vertices in the path from  $z_4$  towards the root  $u$  change. Furthermore, consider a vertex  $w$  in the tree that is an ancestor of at least two endpoints  $z_i, z_j$  of back edges  $b_i, b_j$ . The lowpoint of  $w$  does not change when we remove  $b_i$ . These observations lead us to the following lemma.

**Lemma 21** *In a spine of the recursion tree, the update of lowpoints in the certificate by operation `right_update(C, e)` can be done in  $O(|V_X|)$  total time.*

*Proof:* Take each back edge  $b_i = (z_i, u)$  in the order defined by `choose(C, u)`. Remove  $b_i$  from  $lb(u)$  and  $ab(z_i)$ . Starting from  $z_i$ , consider each vertex  $w$  in the path from  $z_i$  towards the root  $u$ . On vertex  $w$ , we update  $lowpoint(w)$  using the standard procedure: take the endpoint  $y$  of the first edge in  $ab(w)$  (the back edge that goes the nearest to the root of the tree) and choosing the minimum between  $\gamma(y)$  and the lowpoint of each child of  $w$ . When the updated  $lowpoint(w) = \gamma(u)$ , we stop examining the path from  $z_i$  to  $u$  since it implies that the lowpoint of the vertex can not be further reduced (i.e.  $w$  is both an ancestor to both  $z_i$  and  $z_{i+1}$ ).

The total cost of updating the lowpoints when removing all back edges  $b_i$  is  $O(|V_X|)$ : there are  $O(|V_X|)$  tree edges and we do not traverse the same tree edge twice since the process described stops at the first common ancestor of endpoints of back edges  $b_i$  and  $b_{i+1}$ . By contradiction: if a tree edge  $(x, y)$  would be traversed twice when removing back edges  $b_i$  and  $b_{i+1}$ , it would imply that both  $x$  and  $y$  are ancestors of  $z_i$  and  $z_{i+1}$  (as edge  $(x, y)$  is both in the path  $z_i$  to  $u$  and the path  $z_{i+1}$  to  $u$ ) but we stop at the first ancestor of both  $z_i$  and  $z_{i+1}$ .  $\square$

Let us now consider the removal from the certificate of vertices that are no longer in  $B_{u,t}$  as consequence of operation `right_update(C, e)` in a spine of the recursion tree. By removing a back edge  $b_i = (z_i, u)$ , it is possible that a vertex  $w$  previously in  $H_X$  is no longer in the bead string  $B_{u,t}$  (e.g.  $w$  is no longer biconnected to  $u$  and thus there is no simple path  $u \rightsquigarrow w \rightsquigarrow t$ ).

**Lemma 22** *In a spine of the recursion tree, the branches of the DFS that are no longer in  $B_{u,t}$  due to operation `right_update(C, e)` can be removed from the certificate in  $O(|V_X|)$  total time.*

*Proof:* To prune the branches of the DFS tree that are no longer in  $H_X$ , consider again each vertex  $w$  in the path from  $z_i$  towards the root  $u$  and the vertex  $y$ , parent of  $w$ . It is easy to check if  $w$  is an articulation point by checking if the updated  $lowpoint(w) > \gamma(y)$ . If that is the case and  $w$  is not in the leftmost path of the DFS, it implies that  $w \notin B_{u,t}$ , and therefore we cut the subtree of  $w$  and bookkeep it in  $I$  to restore later. Like in the update the lowpoints, we stop examining the path  $z_i$  towards  $u$  in a vertex  $w$  when  $lowpoint(w) = \gamma(u)$  (the lowpoints and biconnected components in the path from  $w$  to  $u$  do not change). When cutting the subtree of  $w$ , note that there are no back edges connecting it to  $B_{u,t}$  ( $w$  is an articulation point) and therefore there are no updates to the lists  $lb$  and  $ab$  of the vertices in  $B_{u,t}$ . Like in the case of updating the lowpoints, we do not traverse the same tree edge twice (we use the same halting criterion).  $\square$

With Lemma 21 and Lemma 22 we finalize the proof of Lemma 16. Fig. 5 shows the changes the bead string  $B_{u,t}$  from Fig. 4 goes through in the corresponding spine of the recursion tree.






Figure 5: Example application of `right_update( $C, e$ )` on a spine of the recursion tree

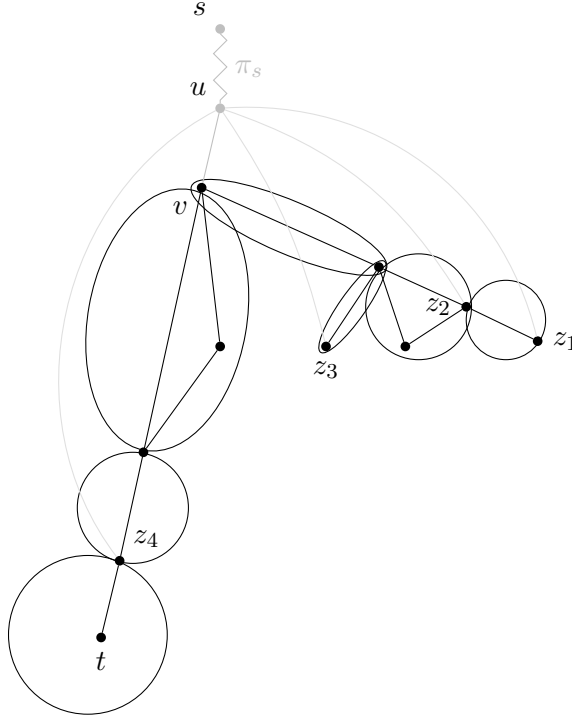


Figure 6: Block tree after removing vertex  $u$

## B.2 Operation $\text{left\_update}(C, e)$ in a spine of the recursion tree

To compute  $\text{left\_update}(C, e)$  in the binary nodes of a spine, we use the fact that in every left branching from that spine, the graph is the same (in a spine we only remove edges incident to  $u$  and on a left branch from the spine we remove the node  $u$ ) and therefore its block tree is also the same. In Fig. 6, we show the resulting block tree of the graph from Fig. 4 after having removed vertex  $u$ . However, the certificates on these left branches are not the same, as they are rooted at different vertices. In the example we must compute the certificates  $C_1 \dots C_4$  corresponding to bead strings  $B_{z_1, t} \dots B_{z_4, t}$ . We do not take into account the cost of the left branch on the last node of spine (corresponding to  $B_{v, t}$ ) as the node is unary and we have shown in Lemma 15 how to maintain the certificate in  $O(1)$  time.

By using the reverse DFS order of the back edges, we are able to traverse each edge in  $H_X$  only an amortized constant number of times in the spine.

**Lemma 23** (Lemma 17 restated) *The calls to operation  $\text{left\_update}(C, e)$  in a spine of the recursion tree can be charged with a time cost of  $O(|E_X|)$  to that spine.*

To achieve this time cost, for each back edge  $b_i = (z_i, u)$ , we compute the certificate corresponding to  $B_{z_i, t}$  based on the certificate of  $B_{z_{i-1}, t}$ . Consider the compacted head  $H_X = (V_X, E_X)$  of the bead string  $B_{u, t}$ . We use  $O(|E_X|)$  time to compute the first certificate  $C_1$  that corresponds to bead string  $B_{z_1, t}$ . Fig. 7 shows bead string  $B_{z_1, t}$  from the example of Fig. 4.

**Lemma 24** *The certificate  $C_1$ , corresponding to bead string  $B_{z_1, t}$ , can be computed in  $O(|E_X|)$  time.*

*Proof:* Let  $t'$  be the last vertex in the path  $u \rightsquigarrow t$  s.t.  $t' \in V_X$ . Since  $t'$  is an articulation point, the

subtree of the DFS tree rooted in  $t'$  is maintained in the case of removal vertex  $u$ . Therefore the only modifications of the DFS tree occur in head  $H_X$  of  $B_{u,t}$ .

To compute  $C_1$ , we remove  $u$  and rebuild the certificate starting from  $z_1$  using the algorithm from Lemma 13 restricted to  $H_X$  and using  $t'$  as target and  $\gamma(t')$  as a baseline to  $\gamma$  (instead of the depth). In particular we do the following. In order to set  $t'$  to be in the leftmost path, we perform a DFS traversal of graph  $H_X$  starting from  $z_1$  and stop when we reach vertex  $t'$ . We then compute the DFS tree, traversing the path  $z_1 \rightsquigarrow t'$  first.

*Update of  $\gamma$ .* For each tree edge  $(v, w)$  in the  $t' \rightsquigarrow z_1$  path, we set  $\gamma(v) = \gamma(w) - 1$ , using  $\gamma(t')$  as a baseline. During the rest of the traversal, when visiting vertex  $v$ , let  $w$  be the parent of  $v$  in the DFS tree. We set  $\gamma(v) = \gamma(w) + 1$ . This maintains the property that  $\gamma(v) > \gamma(w)$  for any  $w$  ancestor of  $v$ .

*Update of lowpoints and pruning of the tree.* Bottom-up in the DFS-tree of  $H_X$ , we compute the lowpoints using the lowpoints of the children. Let  $z$  be the parent of  $v$ . If  $\text{lowpoint}(v) > \text{lowpoint}(z)$  and  $v$  is not in the leftmost path in the DFS, we cut the subtree of  $v$  as it does not belong to  $B_{z_1,t}$ .

*Computing  $lb$  and  $ab$ .* During the traversal, when finding a back edge  $e = (v, w)$ , if  $w$  is a descendant of  $v$  we append  $e$  to both  $lb(v)$  and  $ab(w)$ ; else we append  $e$  to both  $ab(v)$  and  $lb(w)$ . This maintains the DFS order in the back edge lists.

This procedure takes  $O(|E_X|)$  time.  $\square$

To compute each certificate  $C_i$ , corresponding to bead string  $B_{z_i,t}$ , we are able to avoid visiting most of the edges that belong  $B_{z_{i-1},t}$ . Since we take  $z_i$  in reverse DFS order, on the spine of the recursion we visit  $O(|E_X|)$  edges plus a term that can be amortized.

**Lemma 25** *For each back edge  $b_i = (z_i, u)$  with  $i > 1$ , let  $E_{X_i}'$  be the edges in the first bead in common between  $B_{z_i,t}$  and  $B_{z_{i-1},t}$ . The total cost of computing all certificates  $B_{z_i,t}$  in a spine of the recursion tree is:  $O(E_X + \sum_{i>1} E_{X_i}')$ .*

*Proof:* Let us compute the certificate  $C_i$ : this is the certificate of the left branch of the  $i$ th node of the spine where we augment the path with the back edge  $b_i = (z_i, u)$  of  $lb(u)$ .

For the general case of  $C_i$  with  $i > 1$  we also rebuild (part) of the certificate starting from  $z_i$  using the procedure from Lemma 13 but we use information gathered in  $C_{i-1}$  to avoid exploring useless branches of the DFS tree. The key point is that, when we reach the first bead in common to both  $B_{z_i,t}$  and  $B_{z_{i-1},t}$ , we only explore edges internal to this bead. If an edge  $e$  that leaves the bead leads to  $t$ , we can reuse a subtree of  $C_{i-1}$ . If  $e$  does not lead to  $t$ , then it has already been explored (and cut) in  $C_{i-1}$  and there is no need to explore it again since it is going to be discarded.

In detail, we start computing a DFS from  $z_i$  in  $B_{u,t}$  until we reach a node  $t' \in B_{z_{i-1},t}$ . Note that the bead of  $t'$  has one entry point and one exit point in  $C_{i-1}$ . After reaching  $t'$  we proceed with the traversal using only edges already in  $C_{i-1}$ . When arriving at a vertex  $w$  that is not in the same bead of  $t'$ , we stop the traversal. If  $w$  is in a bead towards  $t$ , we reuse the subtree of  $w$  and use  $\gamma(w)$  as a baseline of the numbering  $\gamma$ . Otherwise  $w$  is in a bead towards  $z_{i-1}$  and we cut this branch of the certificate. When all edges in the bead of  $t'$  are traversed, we proceed with visit in the standard way.

Given the order we take  $b_i$ , each bead is not added more than once to a certificate  $C_i$ , therefore the total cost over the spine is  $O(|E_X|)$ . Nevertheless, the internal edges  $E_{X_i}'$  of the first bead in common between  $B_{z_i,t}$  and  $B_{z_{i-1},t}$  are explored for each back edge  $b_i$ .  $\square$


Although the edges in  $E_{X_i}'$  are in a common bead between  $B_{z_i,t}$  and  $B_{z_{i-1},t}$ , these edges must be visited. Since the entry point in the common bead can be different for  $z_i$  and  $z_{i-1}$ , the DFS

tree of that bead can also be different. For an example, consider the case where  $z_i, \dots, z_j$  are all in the same bead after the removal of  $u$ . The bead strings  $B_{z_i,t} \dots B_{z_j,t}$  are the same, but the roots  $z_i, \dots, z_j$  of the certificate are different, so we have to compute the corresponding DFS of the first bead  $|j - i|$  times. Note that this is not the case for the other beads in common: the entry point is always the same.


**Lemma 26** *The cost  $O(E_X + \sum_{i>1} E_{X'_i})$  on a spine of the recursion tree can be amortized to  $O(E_X)$  in that spine.*

*Proof:* We can charge the cost  $O(|E_{X'_i}|)$  of exploring the edges in the first bead in common between  $B_{z_i,t}$  and  $B_{z_{i-1},t}$  to another node in the recursion tree. Since this common bead is the head of at least one certificate in the recursion subtree of the left child of the  $i$ th node of the spine. Specifically, we charge the first and only node in the *leftmost* path of the  $i$ th child of the spine that has exactly the edges  $E_{X'_i}$  as head of its bead string: (i) if  $|E_{X'_i}| \leq 1$  it corresponds to a unary node or a leaf in the recursion tree and therefore we can charge it with  $O(1)$  cost; (ii) otherwise it corresponds to a first node of a spine and therefore we can also charge it with  $O(|E_{X'_i}|)$ . We use this charging scheme when  $i \neq 1$  and the cost is always charged in the leftmost recursion path of  $i$ th node of the spine, consequently we never charge a node in the recursion tree more than once.  $\square$

Lemmas 25 and 26 finalize the proof of Lemma 17. Fig. 7 shows the certificates of bead strings  $B_{z_i,t}$  on the left branches of the spine from Fig. 4.



(a) Certificate of bead string  $B_{z_1,t}$  (b) Certificate of bead string  $B_{z_2,t}$  (c) Certificate of bead string  $B_{z_3,t}$



(d) Certificate of bead string  $B_{z_4,t}$

Figure 7: Certificates of the left branches of a spine